Performance Benefits of VirtualChannel™ Memory

Name : Jeffery H. Lee
Title : Director, Future Technology Enabling
Company : ELPIDA Memory (USA)

San Jose January 23-24, 2001 • Taipei February 14-15, 2001

VirtualChannel™

ELPIDA
The Power of Focus

Evolution of Super High-Bandwidth Memory

Contents

- Consideration of the Memory Bandwidth
- Solution for Multi-Tasking, Low-Latency and High-Bandwidth Memory
- Comparison of DRAM Architectures
- Future VCMemory™

Consideration of the Memory System

Memory Issues

- Big overhead in row switching
- Only one open row in one bank
- Shared pages (open rows) by all memory masters (CPU, AGP, PCI, ...)
- Not fast enough Page Hit Latency
- Insufficient flexibility in scheduling operations

DRAM Access Time Trend

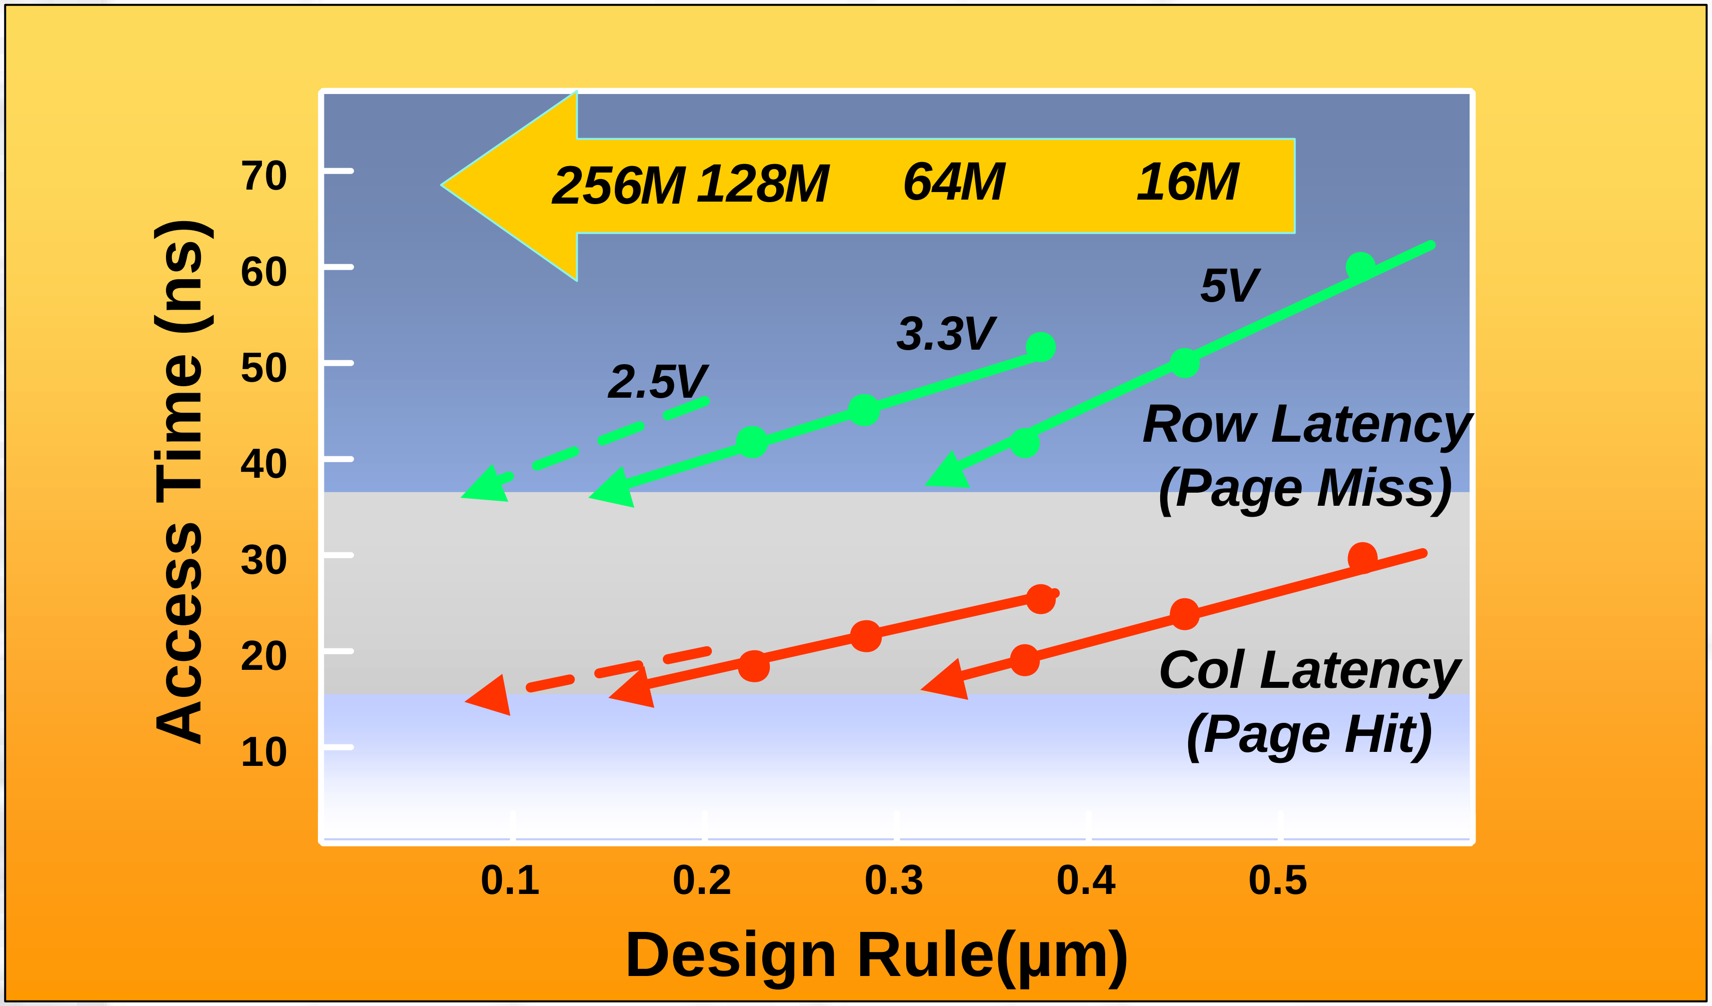

DRAM Cycle Time Trend

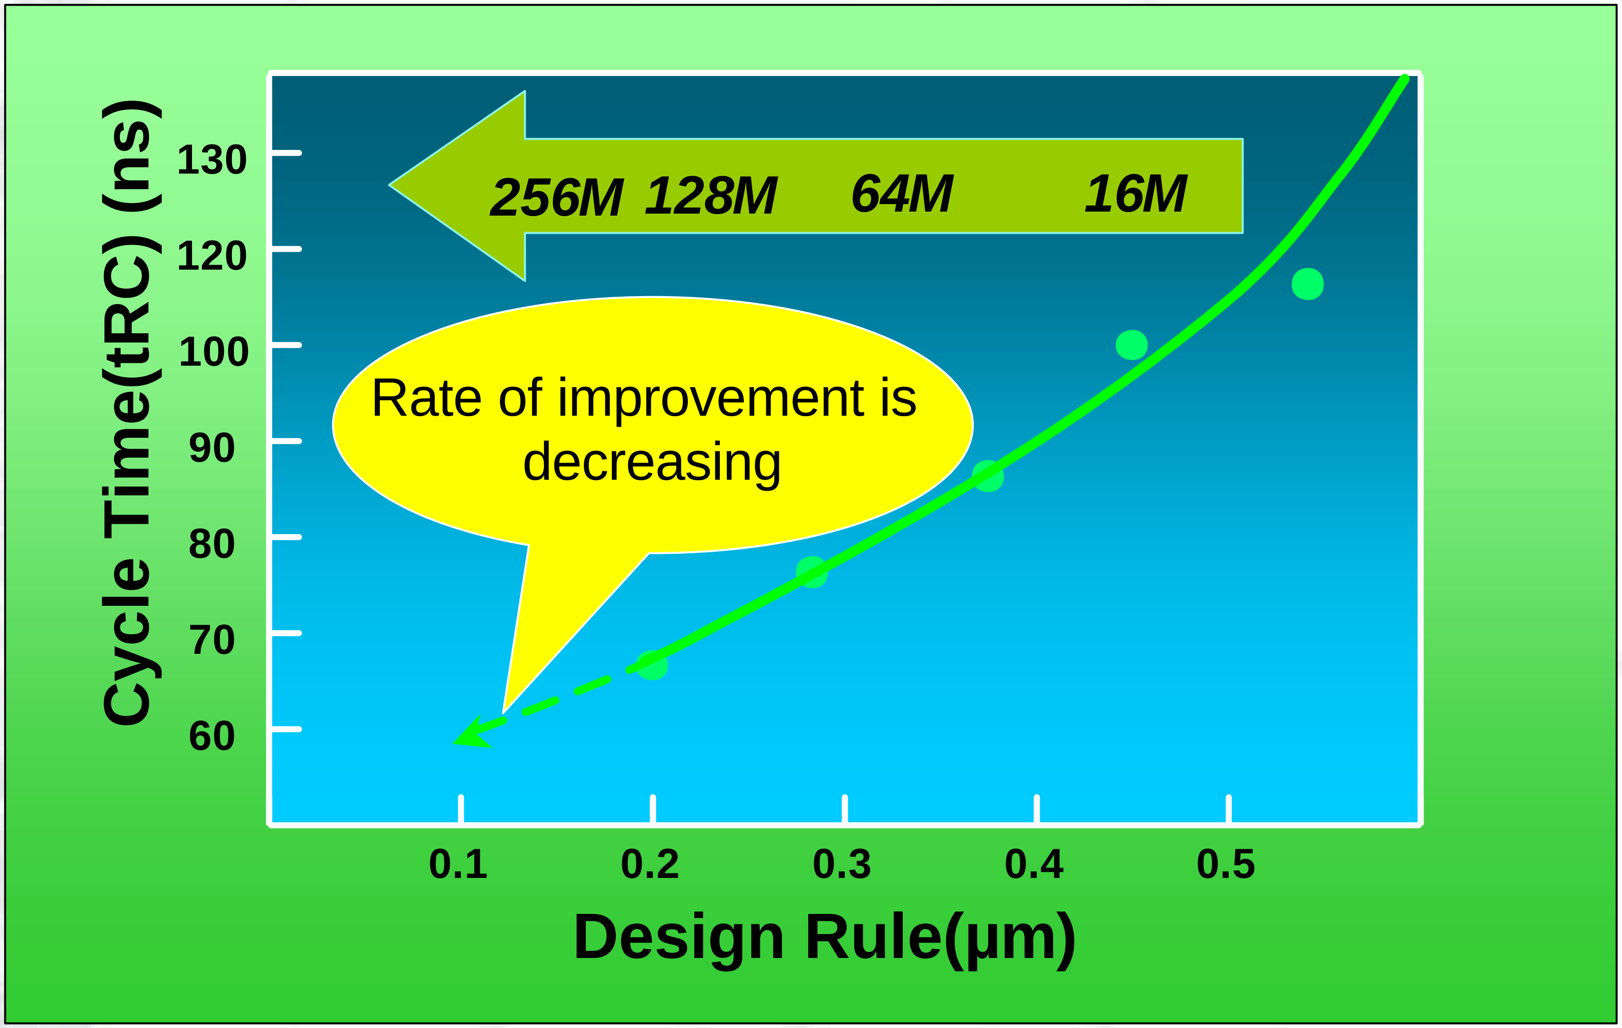

Memory Bandwidth

- Two Bandwidths
  - Memory bus bandwidth
    - Logical Interface : Protocol type, Parallel
    - Physical Interface : LVTTL, SSTL, Rambus
    - Packages : TSOP, BGA, CSP, SVP etc..
  - Memory access bandwidth
    (Sustained rate of read/ write per second)
    - Memory core architecture
    - Multi-Bank, VirtualChannel…
    - Memory core access

Memory Bus Bandwidth

- Memory chip I/O speed
  - DDR266, DDRII and very challenging speed
- Memory bus width
  - System memory bus width
  - Memory module internal bus width
- Reduced bus collisions
  - Equalizing read/ write latency
  - Posted CAS, Delayed write

Memory Access Bandwidth

- **Higher concurrent bank operations**
  - More banks or more channels
- More page hit, less miss/empty page
  - **More open pages**
  - **Less page miss**
- **Lower refresh overhead**
  - **Higher core voltage/ Lower junction temp**
  - **Concurrent bank read/ write and refresh**
  - **Preserving opened page context while refresh**

Memory Access Bandwidth

- **Command bandwidth**
  - Clock rate
  - **Compound operations**
- Low page hit latency
  - **Fast circuits**
  - **Die size and power**
- **Low page miss/ empty latency**
  - **Higher density vs. smaller geometry**
  - **Lowest possible cost**

DRAM Core Timing Projections

Operation type	Timing Parameter	DDRII-600	DDRII-400	DDR266	PC133	PC100	Units
Data Clock period	tdcp	1.67	2.5	3.75	7.5	10	ns
Activate to Activate	trrd	13.36	15	15	15	20	ns
Ras to Cas	trcd	13.36	15	15	22.5	20	ns
Cas data latency	tcl	13.36	15	15	22.5	30	ns
Row Precharge time	trp	13.36	15	15	22.5	20	ns
Row Cycle time	trc	40.08	45	45	67.5	70	ns
Last data in to prechrg	trDL	13.36	15	15	15	20	ns
Last data in to Act	tDAL	26.72	30	30	30	30	ns
Page Empty latency	PGempty	26.72	30	30	45	50	ns
Page Miss latency	PGMissed	40.08	45	45	67.5	70	ns
Refresh Time	RefreshT	53.44	60	60	90	90	ns

DRAM Core Timing Projections

Operation type	Timing Parameter	DDRII-600	DDRII-400	DDR266	PC 133	PC100	Units
Activate to Activate	trrd/tdcp	8	6	4	2	2	tdcp
Ras to Cas	trcd/tdcp	8	6	4	3	2	tdcp
Cas data latency	tcl/tdcp	8	6	4	3	3	tdcp
Row Precharge time	trp/tdcp	8	6	4	3	2	tdcp
Row Cycle time	trc/tdcp	24	18	12	9	7	tdcp
Last data in to prechrg	trDL/tdcp	8	6	4	2	2	tdcp
Last data in to Act	tDAL/tdcp	16	12	8	4	3	tdcp
Page Empty latency	PGempty	16	12	8	6	5	tdcp
Page Miss latency	PGMissed	24	18	12	9	7	tdcp
Refresh Time	RefreshT	32	24	16	12	9	tdcp

DRAM Core Timing Projections (GRAPHICS)

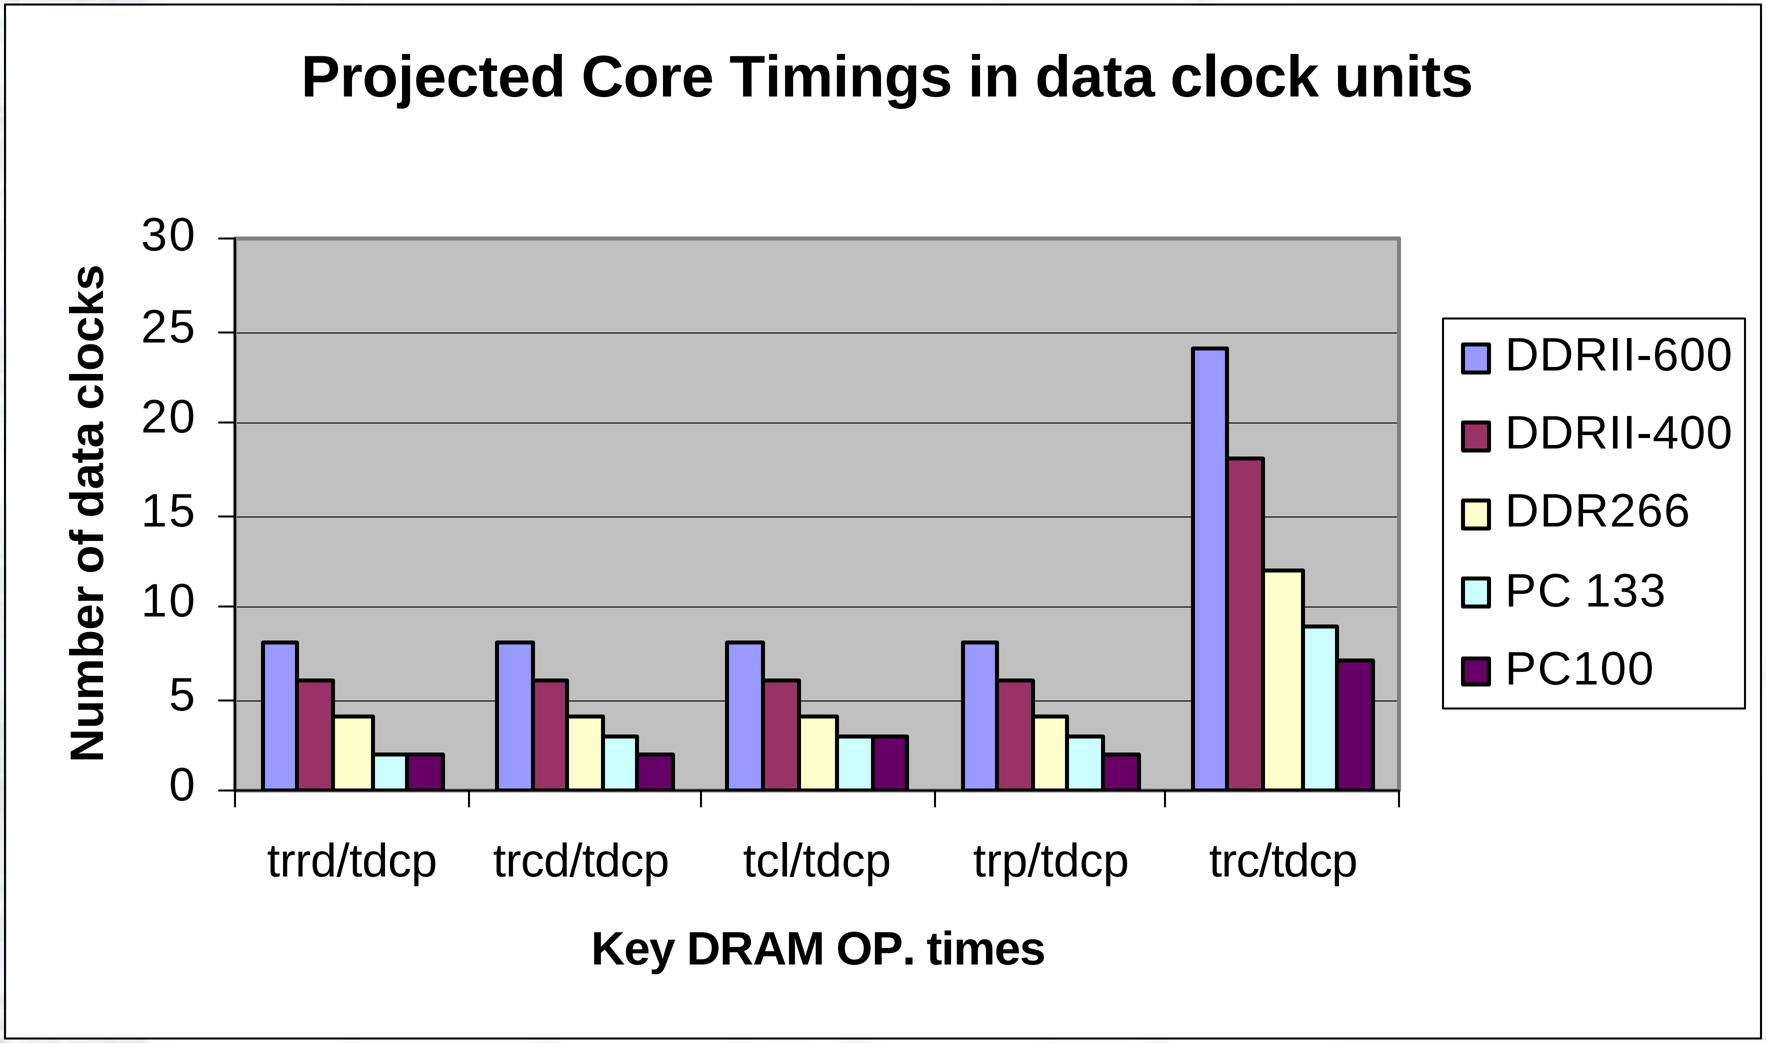

DRAM Core Timing Projections (GRAPHICS)

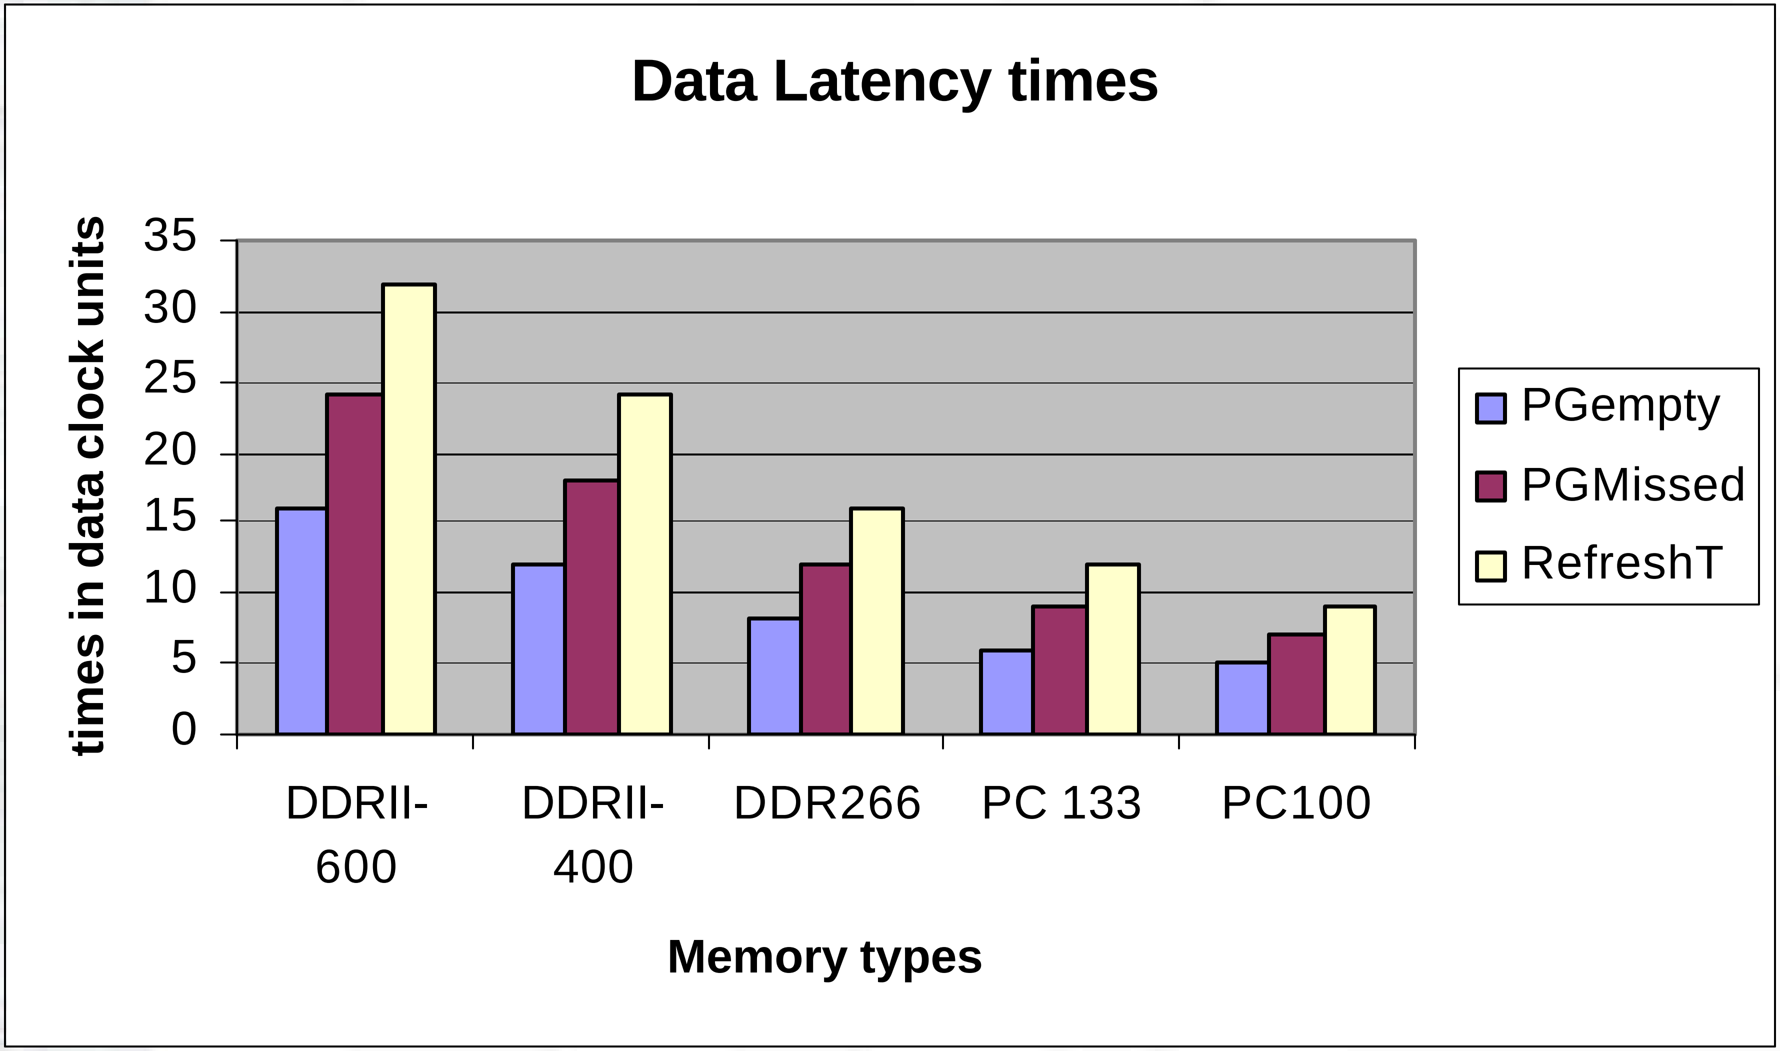

Summary of DRAM issues

- As DRAM moves from PC133 to DDR266 to DDR600
  - Page miss and tRC slows down from 9 clocks to 24 clocks.
  - Refresh goes from 12 clocks to 32 clocks
- Multibanks can only partially mitigate the DRAM core problems.
  - 4 banks to 8 banks to 16 banks : + 4%, + 8% die size.
  - Hit rate goes from 60% to 70% to 80%.
- Lower refresh overhead
  - Higher core voltage / Lower Junction temp.
  - Concurrent bank read/write and refresh
  - Preserving opened page context while refresh

Summary of DRAM issues

- Lower refresh overhead
  - Higher core voltage / Lower Junction temp
  - Concurrent bank read/write and refresh
  - Preserving opened page context while refresh 8
- Refresh frequency
  - 16k in 64ms or 8k in 32ms i.e: 4 us per refresh
  - 8 pages are opened. At refresh, all these pages are closed, and the open page context in the chipset is also destroyed.
  - Command - Prch- Act - Prch : 80 ns. (640 ns)
  - It takes 800 ns to recover the equivalent contexts of these pages.
  - Overhead % : 1.6us / 4us .

Summary of DRAM issues

- 8 Multibanks can only partially mitigate the DRAM core problems.
  - 4 banks to 8 banks to 16 banks : + 4%, + 8% die size.
  - Hit rate goes from 60% to 70% to 80%.
- Lower refresh overhead
  - Higher core voltage / Lower Junction temp.
  - Concurrent bank read/write and refresh
  - Preserving opened page context while refresh

Solution for Multi-tasking, Low-Latency and High-Bandwidth Memory

New Applications trends

- 1- 2 GHz CPU with Speculative prefetches.
    - Pipelined prefetches 7 to 8 deep.
    - Processor access rate of 2 G bytes / sec after L2.
- AGP graphics of 1- GBytes /sec
- Multitaskings of Stream applications.
    - Min. 3 up to 5/6 localities
    - 3 -4 tasks going currently in additions to CPU.

VirtualChannel™ Memory

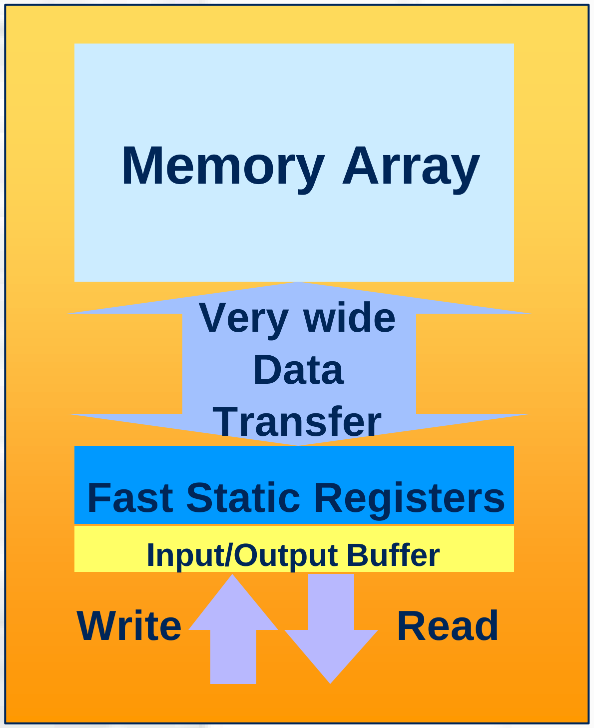

Very Wide Bus **for data transfer for...**

- ***High internal data transfer B.W***
- ***Increasing Bus Efficiency***

Multiple Big Dual-Port Registers **between Array and I/O for…**

- ***Reducing Hit Latency***
- ***Increasing Page Hit Rate***
- ***More Concurrent operation***

VirtualChannel Memory is ...

- ***High Performance Memory Core Architecture for***
  - Very High Bandwidth Memory
  - Any type of memory products (DRAM, ROM, FLASH...)
  - Any type of Interface (LVTTL, SSTL, Rambus®...)
  - All Multi-tasking systems (PC, Server, WS, 3D-Graphics...)

Rambus is registered trademark of Rambus Inc.

Concept of VCMemory

- Very Wide Internal Data Transfer Bus
  - Internal data transfer BW > 10GB/sec (2048bit/20ns)
    => more than 6.4GB/sec I/O BW
- Concurrent operation of Banks and Channels
  - Hide DRAM dead periods behind Channels
- Localities (open rows = Channels) isolated from memory Banks
  - Multiple open rows in one Bank
  - Very fast access to open rows
  - Provide each memory master its own memory resource

Conceptual Schematic

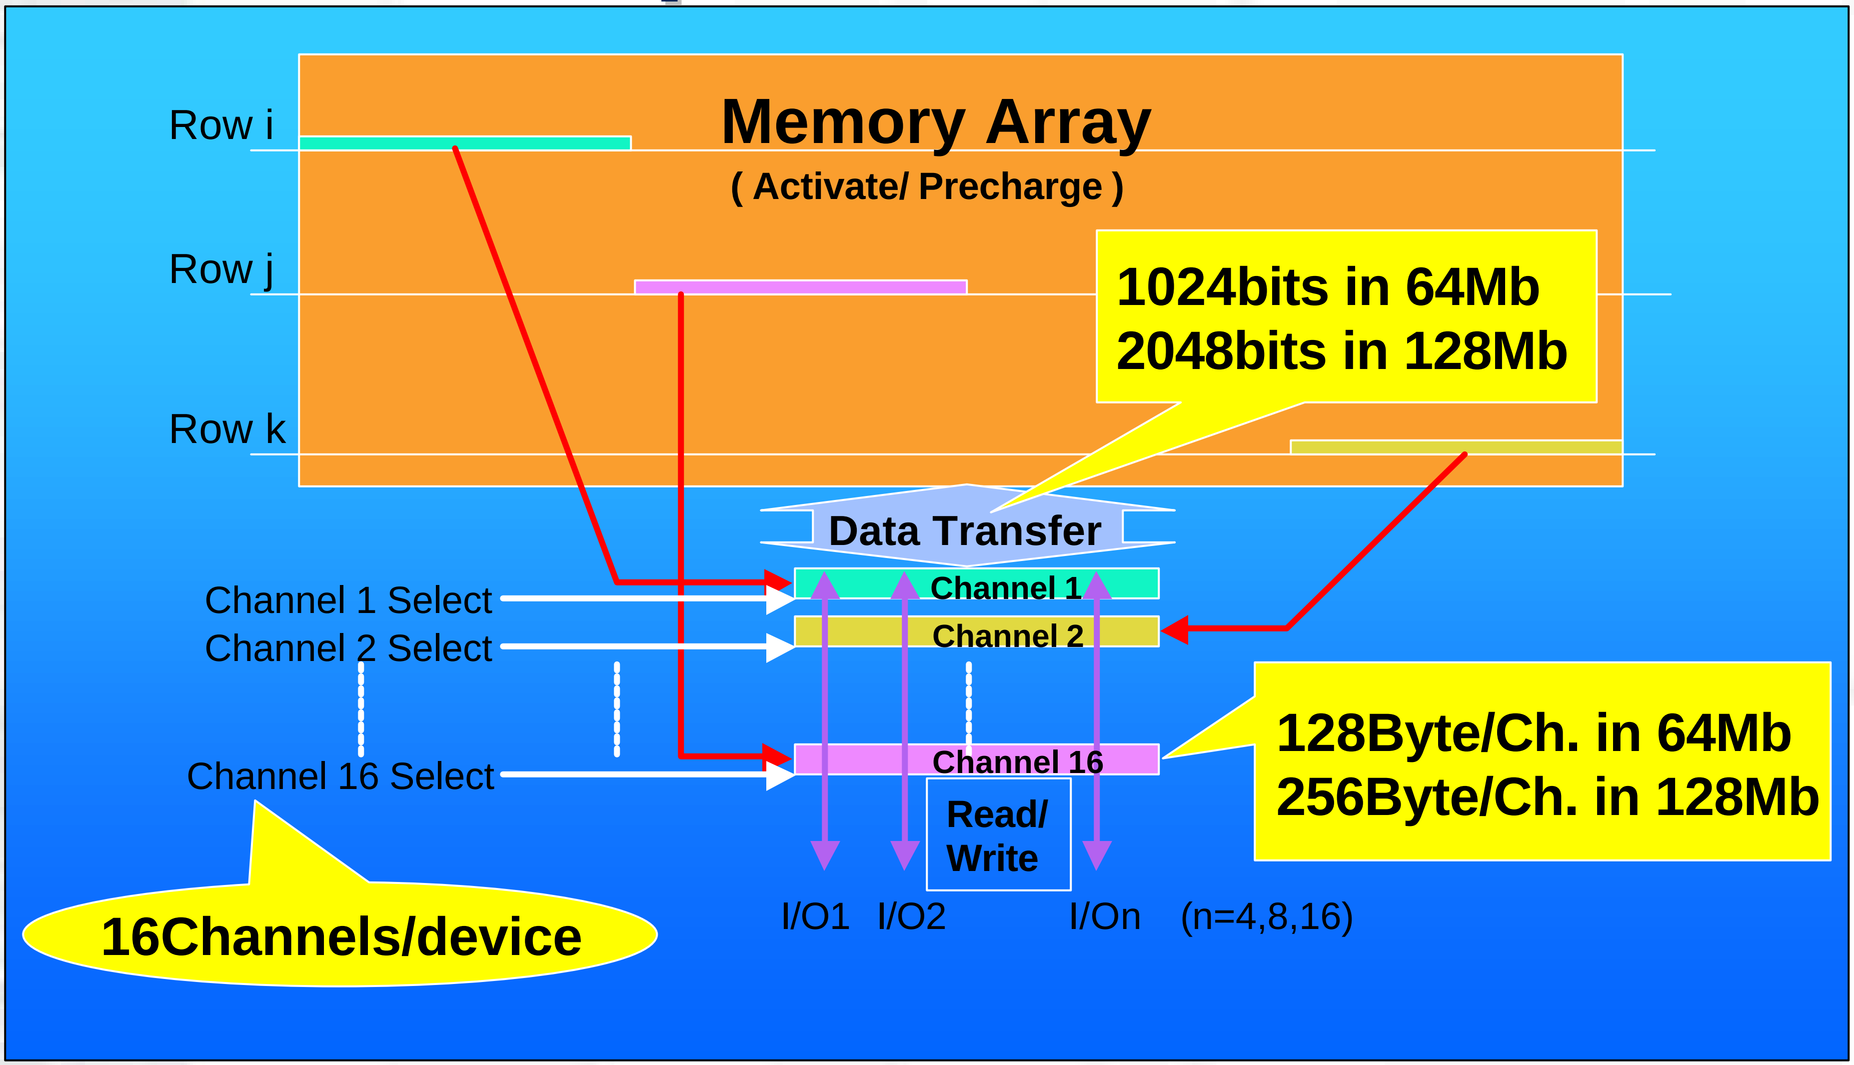

Comparison of DRAM Architectures

DRAM Core Architectures

- **Choices of Memory Core Architecture**
  - Multi-bank, VirtualChannel
- **Each Core Architecture has both strength and weakness**
  - **Performance** (memory access type dependent)
    - latency, bandwidth, effective throughput
  - **power consumption**
  - **cost**
    - die size overhead, wafer fabrication process

Multi-Bank Memory

Strength

- Possible to hide Precharge time (tRP) and RAS to CAS delay (tRCD) by multiple bank interleave.
- Quick access from open page

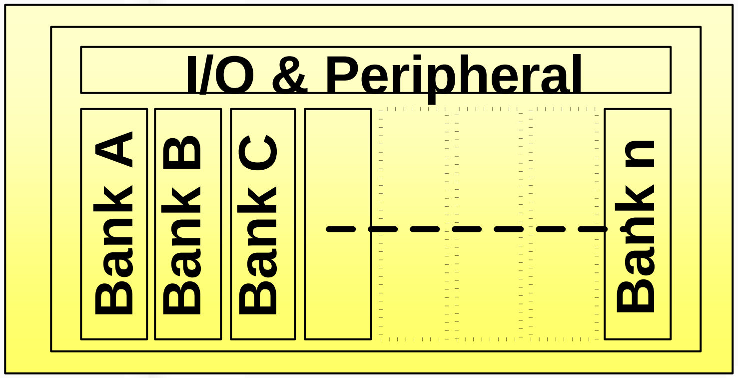

Weakness

- Only one open page/bank
- Big power consumption due to multiple activated banks
  - System performance depends on number of activated bank
- Die size overhead (Decoders, Internal I/O bus…)

VirtualChannel Memory

Strength

- Very high internal BW > 10GB/s
- Low Power consumption
- Easy to hide tRCD, tRP and tRC because of concurrent operation of Banks and Registers(channels)
- Fully associative register architecture

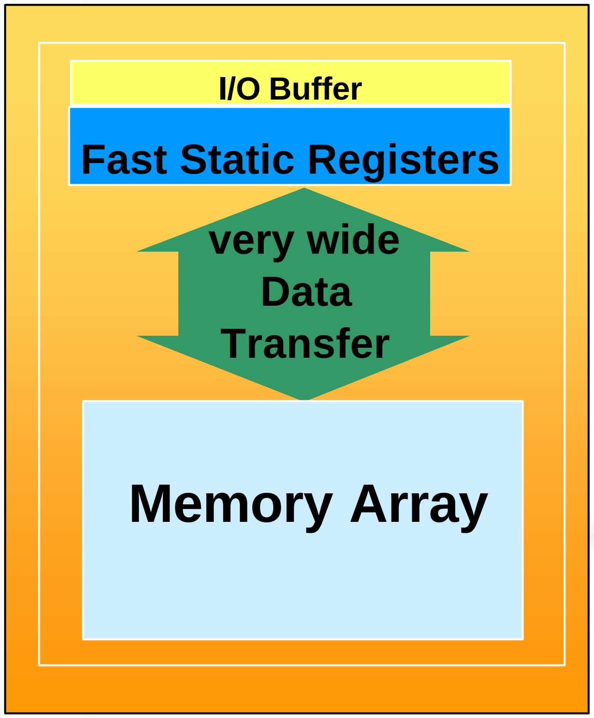

Weakness

- Higher command bus traffic
- Big “Restore” overhead
- Slow “Random Row” cycle time

VirtualChannel^{TM} DRAM System Design

- *Same requirements as for SDRAM*
  - Main board PCB design
    - PC100 and PC133 - supportable data I/O performance
  - Memory module PCB design
  - Memory bus interface
  - Memory package
  - Pin compatible
- *Points requiring modification*
  - Memory controller (Chipset)
    - Necessary of supporting VirtualChannel Memory
  - Internal architecture of the memory (Unit)
    - VirtualChannel DRAM (16 channel buffers)

Why VCMemory is Low Power?

VCMemory dose have

- **Multiple Registers & Very High int. bandwidth**
  - **Higher Page hit rate due to full associative registers**
  - **No need to keep page open mode**
    - **Just “Transfer data and close”**
- **Concurrent operation of Banks and Channels**
  - **Easy hide dead cycles on DRAM, such as tRCD, tRP...**

<u>No need to open multiple banks simultaneously</u>

Data Output (BL=4) (VCDRAM vs. SDRAM)

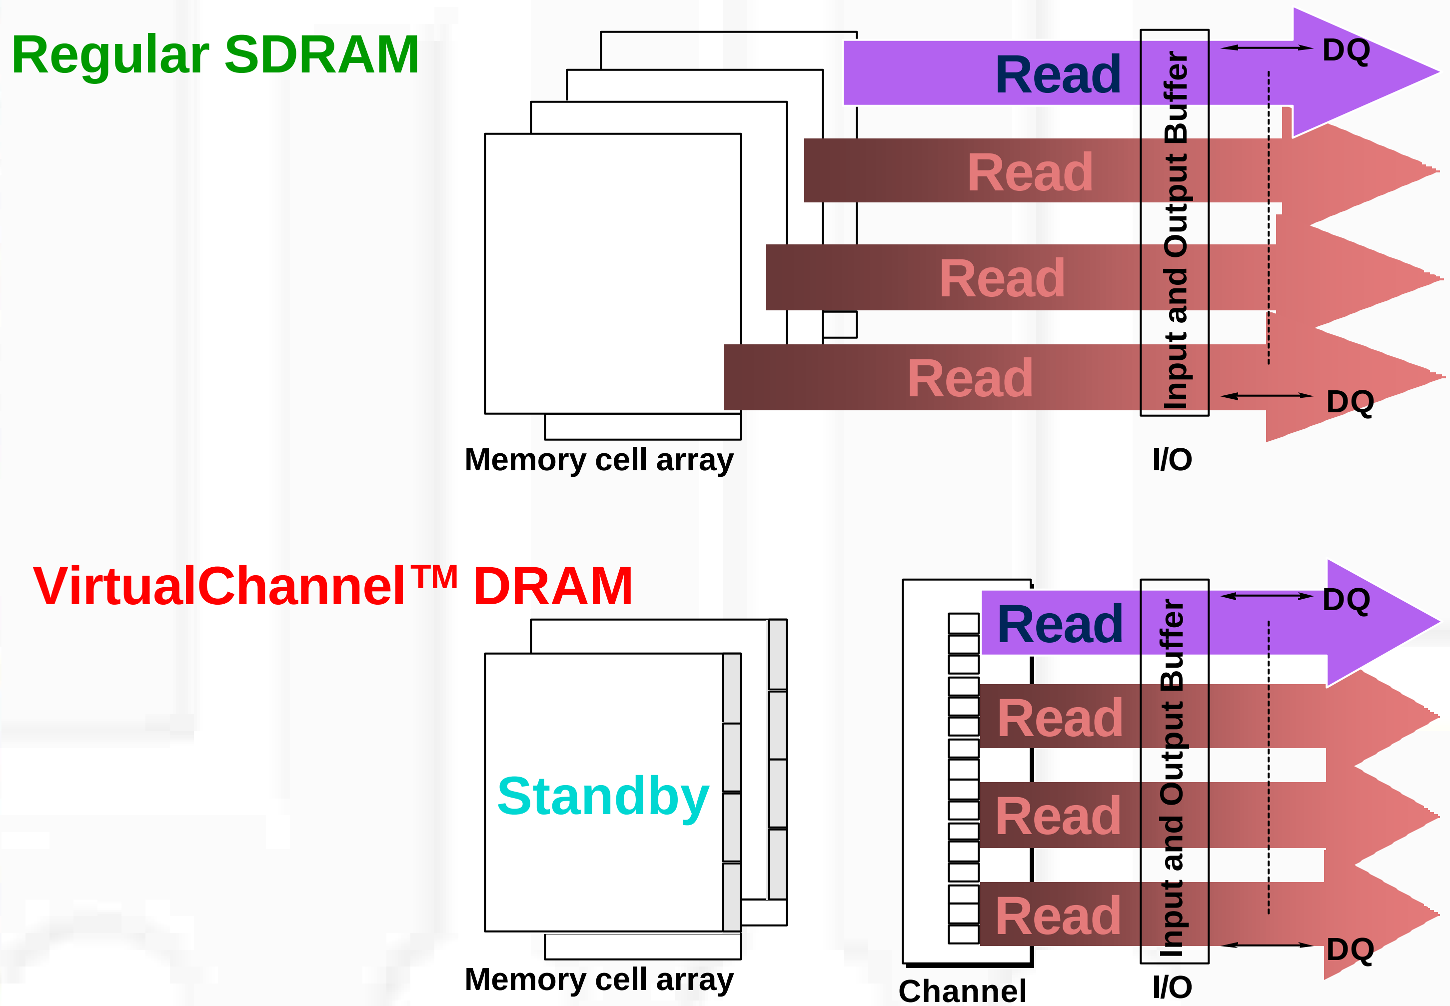

VCMemory™ for Higher BW

- Ready for future requirement
  - BW requirement : 1.6GB/sec =>3.2=>6.4
    - Internal data transfer BW must be higher than I/O BW.
  - Traditional Memory Architecture will be unable to keep up this requirement.
    - Die size overhead increases steeply with internal bus width
    - Many banks helps BW, but need to deal with Power issue
  - VCMemory™ already achieved 10GB/sec

Chip Image

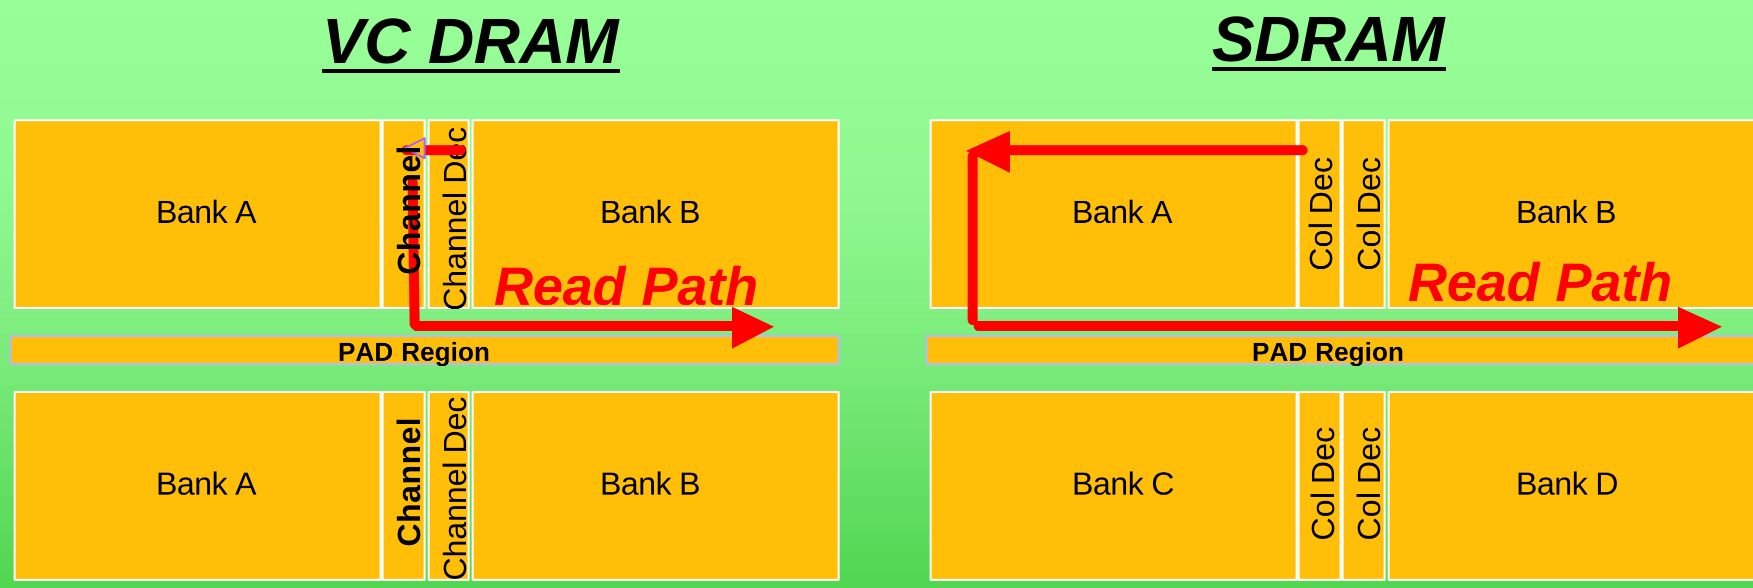

Short Read Path makes Fast Access Time !

The 128Mbits VCDRAM

- VCDRAM architecture is inherently higher frequency and lower latency
- Going from 133MHz to 154MHz
  - 1064 MB/sec 1232 MB/sec (~ 14%)
  - CL = 2 is maintained.
    - 15 ns to 13 ns (13% improvements)
- Die size delta: 5%(64Mb) to 2.5% (128Mb)

The 128Mbits VCDRAM

- Going from 128 byte channel to 256 byte channel
  - 64 bit wide memory goes from 1K bytes per channel to 2K bytes

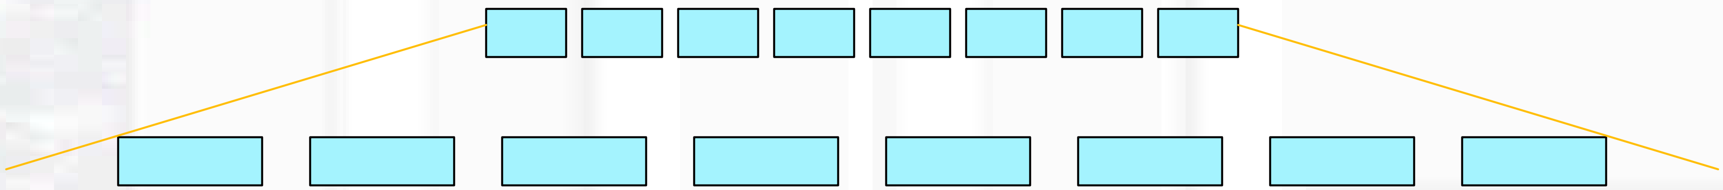

  - Channel Hit rate improves by 15 - 25%
  - Negligible change in Chipset to get full benefits.

Observations/ Issues

- Confusion abounds regarding DRAM performance.
  - What matters? Bandwidth, latency or both?
  - Can DRAM performance be benchmarked?
- Non-DRAM-related factors can cloud the picture.
  - Dependencies include device drivers, chip sets, graphics controllers, hard disk drives, hard disk drive controllers, and others.
  - Performance varies based on board design and BIOS configurations. Not all board makers tune for performance.

A Brief History

VirtualChannel DRAM PC released (NEC) — 1/2001

MP of 128M bits VirtualChannel DRAM available — 5/2000

VirtualChannel DRAM PC released (Micron Electronics) — 9/99

Hyundai support for 64M bits VirtualChannel SDRAM announced — 2/99

Infineon Second Source Announced — 9/98

JEDEC Approval — 8/98

ARM/NEC Joint Announcement — 7/98

6/98 — 64M bits VirtualChannel SDRAM Sampling

1/98 — VIA, ALi, SiS Support Announced

12/97 — JEDEC Proposal

11/97 — Comdex Forum

6/96 — Project Start

NEC Desktop PC with VCDRAM

VC DRAM Roadmap

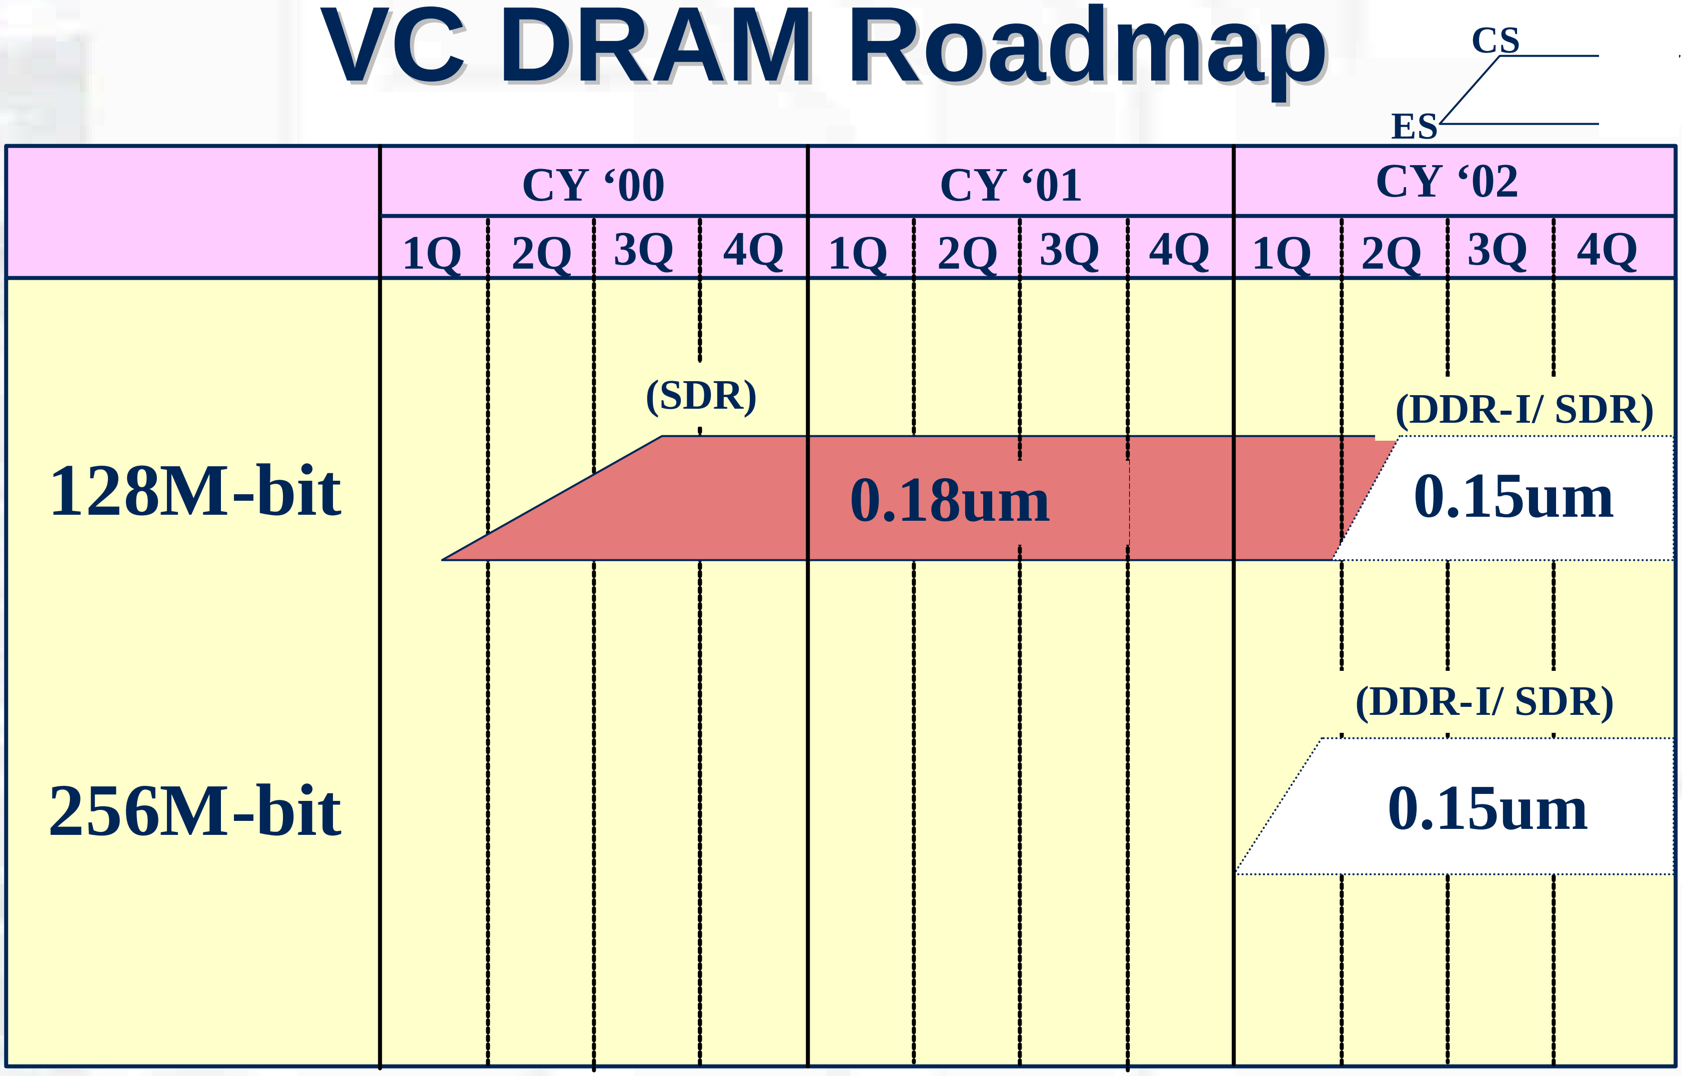

This information may change later.

Future VCMemory

Weakness of VCMemory (SDR SDRAM)

- Higher command bus traffic
- Big “Restore” overhead
  - 160ns to write into memory cells
- Slow “Random Row” cycle time
  - 40ns minimum row switching (interleave) time
- Different latency for write and read
- Insufficient Peak Bandwidth

Ideas for Future VCMemory

- ***Combine commands*** to reduce command bus traffic
- ***Increase throughput***
- ***Increase Number of Banks*** to improve "random row" cycle time

Combining Commands

- Combine multiple commands into one command (especially for background commands)
  - **Example**
    - ACT + Prefetch => ACPR
    - ACT + Prefetch with Auto Precharge => ACPRA
    - RST + ACT(R) => ACRST
    - RSTA + ACT(R) => ACRSA
- Issues of combining commands
  - needs more address bits for channel and segment assignment

Increasing throughput & number of Banks

- Increasing Bandwidth
  - **DDR, QDR, ODR ??**
- **Increasing Banks**
  - **<u>No more than 4 !</u>**
    - Good for power consumption
  - **More than 4 internal banks gives you...**
    - More die size overhead
    - Very small performance gain for VCMemory due to full associative channels.

Benchmark Results

Percentage Performance Advantage

	i840/820	i820/VC	VC/SD	i820/SD
Content Creation	3.9	0.2	3.1	3.3
Content Creation + MP	4.9	2.0	3.4	5.5
Hi. Winstone 99	0.0	-1.0	0.2	-0.8
DualPInsp	0.5	0.2	3.6	3.8
DualPMicroSE	1.1	0.3	2.4	2.7
DualPPhoto	0.0	4.0	5.4	9.6
DualPC++	0.8	-3.6	4.5	0.7
Quake_low	4.2	-2.7	-0.6	-3.4
Quake_high	-0.7	-0.1	-0.1	-0.3

Normalized Benchmark scores

	Intel 840	Intel 820	VC-133	PC-133		
Content Creation	45.1	43.4	43.3	42		
	43.5	42.6	41.7	40.6		
Content Creation + MP	42.5	40.5	39.7	38.4		
	21.6	21.3	19.3	19.4		
Hi. Winstone 99	48.5	48.5	49	48.9		
	51.3	49.1	49.9	50.7		
DualPInsp	31.7	31.55	31.5	30.4	Scaled x5	
	25.95	25.25	25.3	25.2		
DualPMicroSE	57.45	56.85	56.7	55.35	Scaled x15	
	49.2	48.15	48.6	48.6		
DualPPhoto	42	42	40.4	38.32	Scaled x 4	
	32.92	31.68	30.8	30.2		
DualPC++	43.1	42.75	44.35	42.45	Scaled x5	
	33.05	32.05	32.35	32.35		
Quake_low	42.15	40.47	41.61	41.88	Scaled x 0.3	
	40.98	39.93	40.86	41.88		
Quake_high	34.25	34.5	34.55	34.6	Scaled x 0.5	
	41.4	41.25	40.95	40.95		

Anandtech Benchmarks July 2000

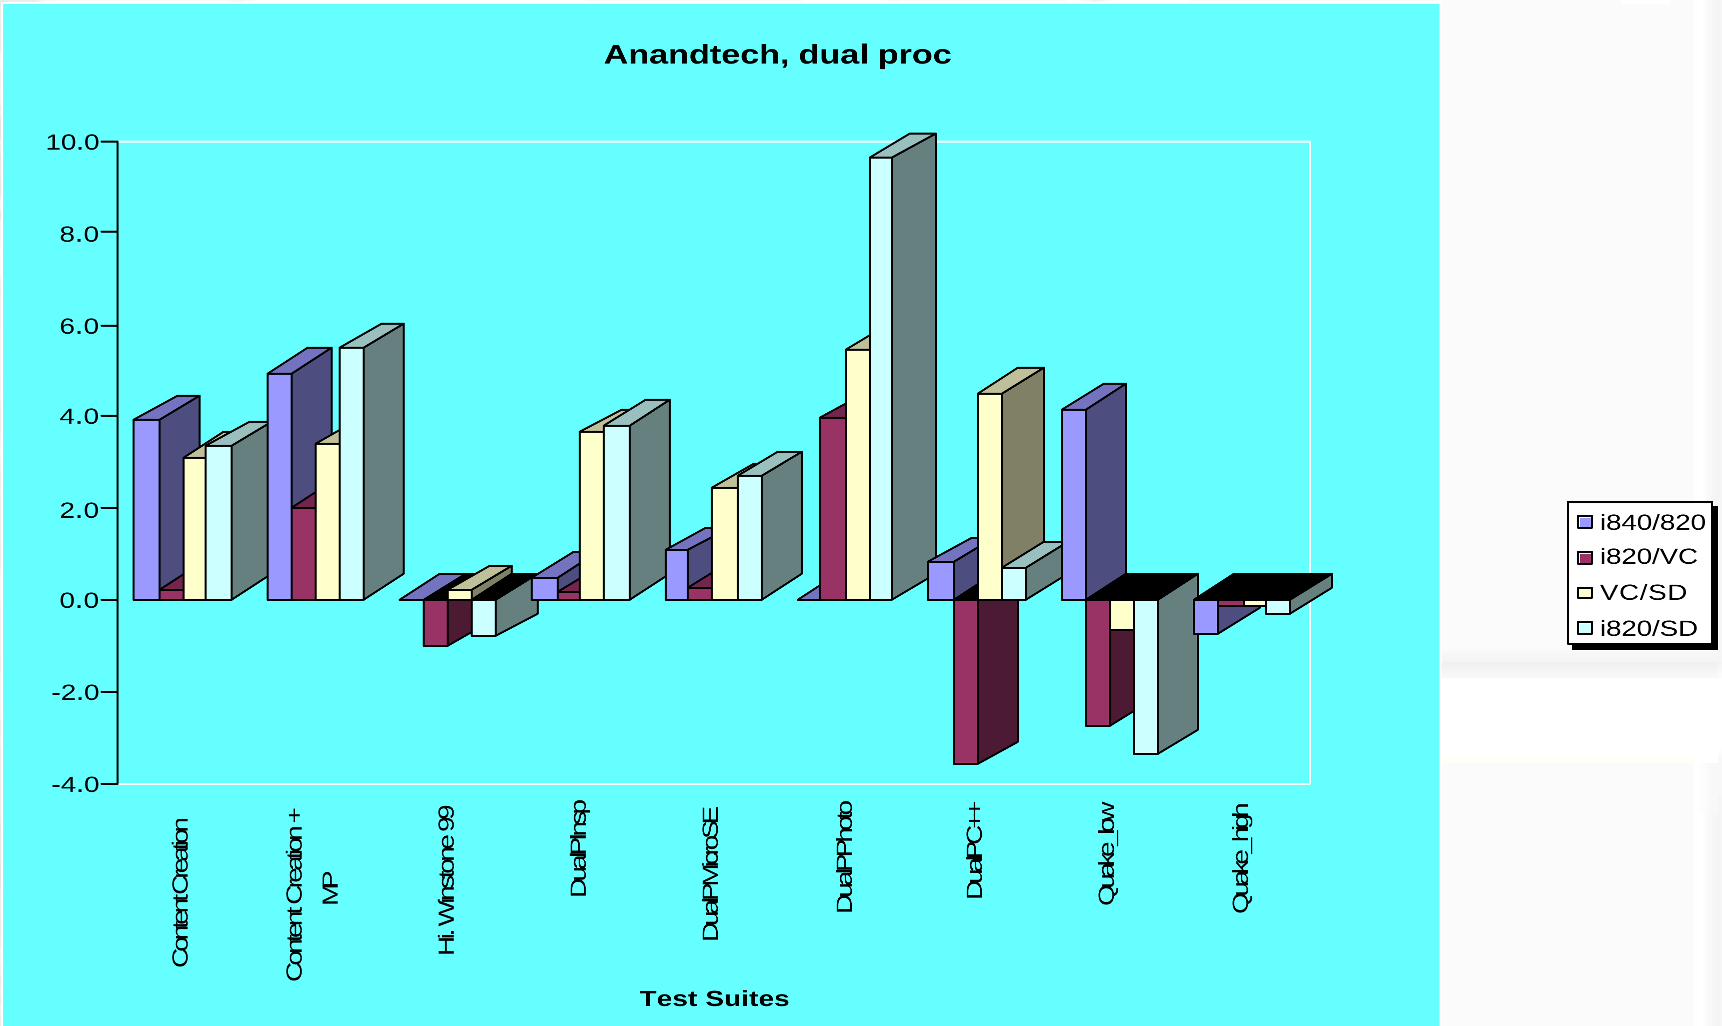

Summary

? VC is a very effective memory core architecture for toadies SDRAM and DDR.

? As DRAM bandwidth increases to satisfy

  ? Higher speed CPUs,

  ? front side bus

  ? multitasking of streaming applications,